SEA-OF-CELLS ARRAY OF TRANSISTORS

BACKGROUND OF THE INVENTION

[01]   The invention concerns an Integrated Circuit (IC) architecture in which individual

5   transistors, each of which resides in a "cell," are arranged in a matrix-like array, thereby

forming a "sea" of the cells.

[02]   Groups of the cells are interconnected among themselves, by local interconnect, into

functional units. (Some of these units are called "MACROS.") The local interconnect in the

units is prohibited from occupying certain layers, such as second-layer metal. The prohibited

10   layer is used instead to connect the individual units to each other.

[03]   Several practices, common in the prior art, tend to utilize resources in integrated

circuits (ICs) in an inefficient manner. These are:

1.   The use of metal level 2 for local interconnect.

2.   The use of metallization located above a row of transistors for interconnect for

15   other transistors, rendering the row of transistors non-usable.

3.   The use of a cell spacing (or "row pitch") in a MACRO which is different

from that of the rest of the array of standard cells into which the MACRO is embedded.

[04]   These practices will be addressed individually.

Metal Level 2 is Used for Local Interconnect

20   CMOS Inverters Generally

[05]   Figure 1 illustrates a common approach to constructing a CMOS inverter. A p-well

and an n-well are constructed in a silicon substrate 3. A gate G, commonly made of

polysilicon, extends across the wells.

[06]   In the p-well, an electric field produced by the gate G generates an n-type channel

25   (not shown) in which electrons flow from a source S to a drain D. In the n-well, this electric

field generates an opposite type of channel, namely, a p-type channel (not shown), in which

holes flow from a source S to a drain D. This electric field modulates the flow of the electrons and holes, and thus modulates the current flowing through the inverter.

[07]   Electric power for the inverter is provided by bus lines Vss and Vdd. These bus lines are generally fabricated in first-layer metal, or METAL 1 in Figure 2. ("POLY" in that Figure refers to polysilicon.)

Trace T is Generally Located in METAL 2

[08]   The two drains D in Figure 1 are connected by an interconnect trace T. The Inventor herein has observed that this trace T is fabricated using second-layer metal, which is labeled METAL 2 in Figure 2. Locating this trace T in METAL 2 presents obstacles to routing other traces, as Figure 3 illustrates.. For example, trace TT cannot take the path shown, because trace T blocks the way. Thus, the freedom of routing of traces such as TT is limited by the local interconnect traces T.

[09]   (Figure 1 has been simplified for ease of illustration. Insulating layers are not shown, and the vias V have been simplified. Figure 4 illustrates a more detailed view. Vias are not pure vertical columns, as in the simplified Figure 1, but, for various technical reasons, take the form shown in Figure 5.)

When Macros are Embedded, the Power Busses Become Disrupted

Standard Cell Arrays Generally

[10]   The cells of a standard cell array typically contain a simple logic function, such as an inverter, a NAND gate, or a D-flip flop. The transistors in these cells are specifically designed for the drive requirements of the particular cell, and spacing of these transistors depends upon such factors as the location of contacts within the cells.

[11]   The spacing, or pitch, between rows of cells is determined by (a) the number of interconnect lines fabricated from METAL 1 (shown in Figure 2) and (b) the cell height. The cell height, in turn, depends upon the transistor configuration within the cells. Figure 6 illustrates these terms.

[12]   The interconnect lines fabricated from METAL 1 are typically laid out by an automated device, or computer program, called a "router," or "auto-router." Different routers

have different algorithms for laying out the lines, so that different routers will produce different interconnect patterns, even though the end result of the connections may be the same.

[13]   Thus, in general, the row pitch is determined by (a) the router used to interconnect the cells in the standard cell array and (b) the height of the individual cells.

Wiring is Primary Consumer of Space

[14]   It is very important to efficiently arrange the wiring in an IC because, in general, the wiring running from transistor-to-transistor consumes more space than the transistors themselves. (The wiring consists of traces fabricated from the METAL layers shown in Figure 2.) Restated, the size of the IC is generally determined by how efficiently the wiring can be routed and compacted, and not by how many transistors the IC contains.

[15]   In a standard cell array, such as that shown in Figure 6, when more wiring is needed, it is common to use the METAL, shown in Figure 2, which is located between rows of cells in Figure 6, such as at location Ll. If additional METAL 2 is required, the cells are then spaced apart, as shown in Figure 6, so that metal lines can be run between them, as indicated.

Embedding MACROs into Standard Cell Arrays Wastes Space

[16]   MACROS are frequently incorporated into ICs containing standard cell arrays. A MACRO is a block of transistors which have been optimized to perform a specific function. In a MACRO, the layout of the individual transistors, their operating characteristics, and their interconnections may have all been matched to each other for optimum performance. Thus, typically, a MACRO is constructed from different sizes of transistors, which are embedded into the standard cell array as shown in Figure 7.

[17]   Since, in general, the ROW PITCH of the MACRO is different from that of the standard cell array, the power busses Vdd and Vss will be interrupted. To accommodate this interruption, the power busses are re-designed as a ring which surrounds the MACRO. Figure 7 shows such a ring generically.

3

Recapitulation

[18]   Therefore,

1.   The use of METAL 2 for local interconnect presents obstacles to the free routing of other interconnects over the cell, as illustrated in Figure 3.

2.   In a standard cell array, the ROW PITCH is determined by the cell height and the number of lines of METAL 1 interconnect placed between the cell rows by the auto-router.

3.   The row pitch in a MACRO is generally different from that of a standard cell array into which the MACRO is embedded. This different row pitch disrupts the power bus system, requiring a ring of power busses to be formed around the MACRO. This approach wastes space within the IC.

OBJECTS OF THE INVENTION

[19]   It is an object of the invention to provide an improved approach to the layout of an integrated circuit.

[20]   It is a further object to provide a method of compacting interconnections in integrated circuits.

[21]   It is yet a further object of the invention to provide a method for improving the performance of the integrated circuit after the layout has been completed, without requiring a new layout to be generated.

SUMMARY OF THE INVENTION

[22]   In one form of the invention, a MACRO, when embedded within a standard cell array, uses the same row pitch as that of the standard cell array itself.

[23]   In another form of the invention, the interconnect within the standard cell is confined to METAL 1 and polysilicon layers, so that METAL 2 is free for routing over the cell.

[24]   In still another form of the invention, the diffusion layer of the transistors within the standard cells is designed for optimum performance after the layout has been completed, rather than at another time.

BRIEF DESCRIPTION OF THE DRAWINGS

[25]   Figure 1 illustrates a CMOS inverter of the prior art.

[26]   Figure 2 illustrates the different layers of metallization used in IC fabrication.

[27]   Figure 3 illustrates the proliferation of traces T, shown in Figure 1, which can occur when trace T is fabricated in a layer of metallization, such as METAL 2 in Figure 2.

[28]   Figure 4 illustrates the CMOS inverter of Figure 1, but in greater detail.

[29]   Figure 5 is a highly simplified depiction of a via of Figure 4. Like-numbered structures correspond in both Figures. In an actual via, the metal in the via C is a portion of element D1. The outline C in Figure 5 is the cutout in the dielectric layer. Metal layer D1 flows down into the via C.

[30]   Figure 6 illustrates a standard cell array.

[31]   Figure 7 illustrates a common approach to embedding a MACRO in a standard cell array.

[32]   Figure 8 illustrates one form of the invention.

[33]   Figure 9 illustrates an imaginary grid on which traces are laid out in the prior art.

[34]   Figure 10A, 10B, 10C, and 10D illustrate the different effective widths of different traces.

[35]   Figure 11 illustrates different possible ways to pack different traces, of different effective widths, between the traces T of Figures 1 and 3.

[36]   Figure 12 illustrates how different traces of different effective widths can be packed on a layer of metallization which lacks the traces T of Figure 3.

[37]   Figure 13 illustrates embedding of a MACRO into a standard cell array, according to the invention.

[38]   Figure 14 illustrates the WIDTH of the CHANNEL of a field-effect transistor.

[39]   Figure 15 illustrates an analog switch of the prior art.

5

DETAILED DESCRIPTION OF THE INVENTION

Fabricate Drain-Drain Local Interconnect in Polysilicon

[40]   Figure 8 shows an inverter, which performs the inversion function, as does the inverter of Figure 1. However, in Figure 8, the connection C between the drains D is fabricated from the same polysilicon layer as the gate G. Trace C and the gate G are coplanar. The polysilicon layer is labeled POLY in Figure 2.

[41]   Using polysilicon for this interconnect eliminates the trace T in Figure 1 which occupies METAL 2. (Figure 3 illustrates a proliferation of these traces T on METAL 2.) With trace T eliminated, the entire layer of METAL 2 can now be used for routing traces which interconnect MACROS and other functional blocks.. This new availability of METAL 2 provides a significant increase in area available for traces, for two major reasons, as will now be explained.

[42]   First, the metal traces on a given layer are generally parallel, as shown in Figure 9. Traces on different layers are connected by vias, indicated by the dashed lines. It can be shown, based on simple assumptions, that, after clearing of METAL 2 of traces T in Figure 3, the number of metal traces which can be fabricated on METAL 2 is thereby increased by about thirty percent. One contributor to this thirty percent value is the fact that the space formerly occupied by traces T is now available for use, whereas previously it was not.

[43]   Second, an additional benefit will be explained by first explaining a situation common in the prior art. Different traces can possess different effective widths. For example, in Figure 10, trace T1, which contains no associated pads for via connections, has an effective width equal to its own width W1. In this example, W1 is assumed to be 1.0 micron, as indicated. (However, in reality, W1 is typically 0.5 - 2.0 microns.)

[44]   In contrast, trace T2 does contain via-pads P. Its effective width is larger, and equal to the via-pad width W2. In this example, W2 is assumed to be 2.0 microns, as indicated. (However, in reality, W2 is typically 1.0 - 4.0 microns.)

[45]   A minimum separation D in Figure 10B must exist between adjacent traces. The minimum separation D is based on effective widths, determined by an edge such as E1, and

not by edge E2. A distance D of 1.0 microns will be assumed. (In reality, a spacing D of 0.5 - 2.0 microns, for traces of 0.5 - 2.0 microns width, is common.)

[46] This spacing D can be allocated to each trace, by attributing one-half to each side of a trace, as shown in Figure 10C. The one-half spacing is labeled D/2. Consequently, trace T1, effectively becomes a rectangle which is 2.0 microns wide, as indicated in Figure 10C; similarly, trace T2 effectively becomes a rectangle 3.0 microns wide.

[47] One must now inquire how many of these rectangles can be packed between the prior-art traces T in Figure 3.

[48] Figure 11 illustrates the eleven possible ways to pack traces T2 (of width W2 in Figure 11) and T1 (of width W1) between traces T. Distance D1 is assumed to be 9.0 microns. The hatched areas represent unused space.

[49] If one assumes that each of the eleven possible combinations is equally likely to occur in an actual integrated circuit, then the average wasted space is the numerical average of the hatched areas. A statistical analysis of a given IC layout can be undertaken to ascertain the actual probabilities of each combination, and may produce a different answer than a simple numerical average of the possibilities in Figure 11. Nevertheless, it is reasonable to assume that the combinations which leave no wasted space, namely, W2-W2-W2 and W2-W1-W1-W1, will never occur with 100 percent frequency; some wasted space will always occur.

[50] Figure 12 illustrates how the traces can be packed onto METAL 2 layer shown in Figure 2, when the traces T in Figure 3 have been eliminated. The traces T1 and T2 can be packed with maximum density. The only space which is left over is the hatched area at the right. It is clear, based on simple observation, that the wasted space in Figure 12 is less than that in Figure 11.

[51] Further, based on reasonable assumptions, it can be calculated that the total wasted space in Figure 12 will be about 40 percent of the wasted space in Figure 11.

[52] Therefore, by fabricating the trace T in Figure 8 from polysilicon, in the same polysilicon layer as the gate G, the prior art array of traces T in METAL 2 in Figure 3 has been eliminated. This elimination provides a thirty percent increase in space, due to factors which include the recovery of the space occupied by the traces T themselves. This

7

elimination further provides a forty percent reduction in waste, by allowing packing of the type shown in Figure 12 to be attained, as opposed to that of Figure 11.

If Uniform Spacing is Required in Prior Art, Waste is Even Greater

[53]   The eleven possibilities shown in Figure 11 may not all be available. For example, some routing techniques, both computer-executed and manual, follow a rule which requires that all traces lie upon imaginary lines of fixed, uniform spacing, such as the lines H and V in Figure 9. Under such a rule, several possibilities in Figure 11 would not be available. For example, any possibility which mixes W1 with W2 would be unavailable, (because spacing is not equal with this mixing).

[54]   Further, other rules are even more restrictive, and require that the traces be laid out on a grid having the largest spacing, which would equal W2 in Figure 10A. Under this rule, only three possibilities of Figure 11 would be available, namely,

(1)   W2;

(2)   W2-W2; or

(3)   W2-W2-W2.

[55]   The other 8 possibilities are not available.

[56]   Under either of these two rules, some of the possibilities of Figure 11 are eliminated. It is clear that, under the added restrictions, the total wasted space becomes greater than if these possibilities were allowed.

Use Same Row Pitch For MACRO Embedded in Standard Cell Array

[57]   In a second aspect of the invention, a MACRO is embedded within a standard cell array using the same ROW PITCH in the MACRO as in the array, as indicated in Figure 13. The different sizes indicate that the MACRO's cells need not be the same size as the standard cell array's, although they can be.

[58]   Under this layout, the same power busses, Vdd and Vss, feed both the MACRO cells and the standard cell array cells. Further, preferably, within the MACRO, no local interconnect is made on metal level 2 (labeled METAL 2 in Figure 2). That is; within the

MACRO, there are no traces such as T shown in Figure 3 located on metal level 2. Stated again, metal level 2 within the MACRO is preferably free of local interconnect.

[59]   This absence of local interconnect allows great flexibility in routing chip interconnect across the MACRO and also allows compaction of the interconnect into very small spaces which are located over the MACRO.

Optimize Channels in FETs After Layout

[60]   The transistors used contain channels, through which carriers flow from source to drain. One channel is shown in Figure 14. It is possible to fabricate channels having different WIDTHs, by adjusting the masking used in fabrication. Different WIDTHs result in different associated capacitances, such as the gate capacitance, and also result in different transconductances.

[61]   Different capacitances can be desired for different transistors. For example, the channel in a transistor used as a load preferably has a narrow WIDTH, to minimize capacitance. Conversely, the channel in a transistor used as a driver (of a load) preferably has a wider WIDTH, to maximize transconductance.

[62]   Under the invention, a given integrated circuit is laid out, using the above principles, including (a) the clearing of METAL 2 of local interconnect, and (b) the use of the same ROW PITCH in a MACRO and a standard cell array in which the MACRO is embedded. Then, the capacitances of selected transistors are computed. (The capacitances cannot be calculated before this time, because the capacitance of each transistor depends on the traces leading to the transistor, as well as on the traces running next to, and over, it.)

[63]   If this computation indicates that some capacitances are not optimal, then the WIDTHs of the channels of the non-optimal transistors can be changed. The WIDTHs of load transistors can be reduced, and the WIDTHS of driver transistors can be widened.

[64]   The inventor notes that the trimming is probably limited by a factor of about three. That is, if the smallest WIDTH possible is used initially, then it is probable that the WIDTH cannot be increased by greater than a factor of three. Conversely, if the largest possible WIDTH is used initially, then it is probable that the WIDTH cannot be reduced by greater than a factor of 1/3. The primary reason is that the maximum WIDTH must lie within the

9

cell boundaries (cells are shown in Figure 3), and the minimum WIDTH is limited by the minimum masking dimensions available.

[65]   Therefore, under the invention,

(a)   METAL 2 is freed of local interconnect, partly by constructing drain-drain interconnect within the polysilicon layer (POLY in Figure 2).

(b)   MACROS have no local interconnect within METAL 2.

(c)   When MACROS are embedded within standard cell arrays, they are given the same ROW PITCH as the standard cell array itself.

(d)   After layout (which is a conceptual step, done either on paper or by computer) and before fabrication, the relevant capacitances of selected (or all) transistors is computed. If any capacitances are non-optimal, the channel WIDTH is reduced or increased, as appropriate.

Invention Allows Non-Uniform Spacing of Interconnect

[66]   Typical Prior-Art interconnect traces are positioned on an x-y grid of fixed spacing. That is, as shown in Figure 9, all horizontal traces must run along horizontal dashed lines H, and all vertical traces must run along vertical dashed lines V. (As a consequence, since all vias must lie upon an intersection point of a horizontal trace and a vertical trace, the vias will lie on a grid point GP.)

[67]   In contrast, under the invention, the spacing of the traces can be non-uniform. Figure 12 illustrates the non-uniform spacing.

Additional Considerations

[68]   1.   "Local Interconnect" refers to signal traces which run between two locations in a given functional block of transistors, or other components. As an example, Figure 15 illustrates an analog switch. (The discussion above has presumed standard cell arrays of CMOS devices; Figure 15 shows BJTs. The appearance of BJTs is purely incidental.) All of the interconnections in that Figure are "local interconnects," with two exceptions:

(1)   The power and ground lines are not considered local interconnect, because the term is not applied to such lines.

(2)   The line carrying the CONTROL VOLTAGE is not a local

interconnect, because it carries the actuation signal for the analog switch. In general, this actuation signal will originate at a location which is unrelated to, and independent of, the location of the analog switch.

[69]   Therefore, Figure 15 illustrates one definition of "local interconnect." Local interconnect includes traces which carry signals from one component to another within a given functional block, and which assist in the execution of the block's function. The line CONTROL VOLTAGE, which carries the incoming signal, is not local interconnect, because it initiates execution of block's function.

[70]   Another definition of local interconnect is again related to functional blocks. Integrated circuit are frequently constructed using MACROs and other functional blocks contained in a library. Each MACRO and functional block can operate isolated, by itself (provided it receives power and input signals). The traces contained within each isolated MACRO and functional block (except power, input and output traces) are local interconnect.

[71]   2.   The invention applies to integrated circuits having high levels of integration. For example, the invention applies to ICs of overall dimension of 5 x 5 mm, or greater. Such ICs are fabricated using VLSI/ULLSI techniques.

[72]   3.   The discussion above considered parallel power busses, such as Vdd and Vss in Figure 6. In the present context, "parallel" does not mean concentric. For example, if a second ring (for Vss) in Figure 7 were to be fabricated, parallel with the Vdd ring shown, the two rings would not be considered "concentric".

[73]   4.   In Figure 13, the power busses Vdd and Vss run parallel, and they are aligned to both the cells of the standard cell array and the cells of the MACRO.

[74]   5.   Under the invention, a MACRO having a fixed, predetermined layout is embedded in a standard cell array. Since the placement does not alter the design of the MACRO, the MACRO's timing is not be affected by this embedding.

[75] 6. The final gate width adjustments (or channel width adjustments), discussed above, are accomplished by the use of computer simulations. A series of timing simulations of the circuit are run, in which transistor gate widths are incrementally changed, and the resulting change in the logic timing are observed. If timing is improved, the new increment is likely to be retained. If the timing worsens, the new increment is likely to be rejected.

[76] The simulations are repeated until the widths converge on fixed values. This technique is known in the art as simulated annealing.

[77] 7. A third and fourth layer of metal interconnect can be utilized in addition to the interconnect, discussed above, provided by (a) the polysilicon layer, (b) METAL 1 and (c) METAL 2 layers. These additional layers maximize the cell density in the layout by providing layers of metal interconnect that can be freely routed over the cells. That is, these layers are not subject to the restrictions discussed above.

[78] Numerous substitutions and modifications can be undertaken without departing from the true spirit and scope of the invention. What is desired to be secured by Letters Patent is the Invention as defined in the following claims.